



PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 :

G06F 15/16, 13/00

A1

(11) International Publication Number:

WO 00/58856

(43) International Publication Date:

5 October 2000 (05.10.00)

(21) International Application Number: PCT/US00/08410

(22) International Filing Date: 30 March 2000 (30.03.00)

(30) Priority Data:
60/127,116 31 March 1999 (31.03.99) US
09/363,670 29 July 1999 (29.07.99) US(71) Applicant: DIVA SYSTEMS CORPORATION [US/US]; 800
Saginaw Drive, Redwood City, CA 94063 (US).(72) Inventors: TAYLOR, Clement, G.; 215 Hampshire Drive,
Plainsboro, NJ 08536 (US). CHIN, Danny; 4 Strathmore
Place, Princeton Junction, NJ 08550 (US). LERMAN, Jesse,
S.; 141 Finnegan Lane, Kendall Park, NJ 08016 (US).
ZACK, Steven; 1308 Oxmead Road, Burlington, NJ 08016
(US). ASHLEY, William; 1155-C LaRoche Terrace,
Sunnyvale, CA 94089 (US).(74) Agents: MOSER, Raymond, R. et al.; Thomason Moser &
Patterson, LLP, 2-40 Bridge Avenue, P.O. Box 8160, Red
Bank, NJ 07701 (US).(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG,
BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE,
ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP,
KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA,
MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU,
SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG,
UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS,
MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ,
BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE,
CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC,
NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA,
GN, GW, ML, MR, NE, SN, TD, TG).

Published

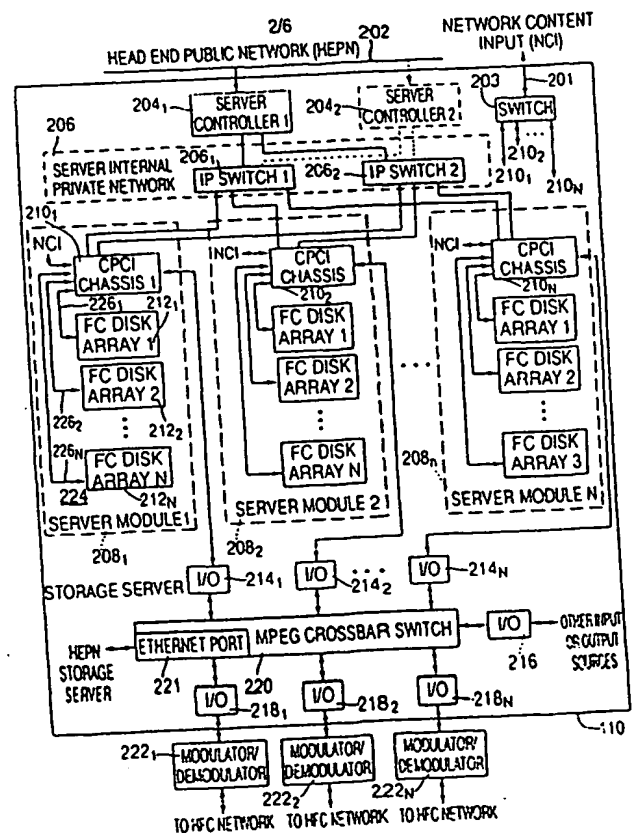
With international search report.

Before the expiration of the time limit for amending the
claims and to be republished in the event of the receipt of
amendments.

(54) Title: TIGHTLY-COUPLED DISK-TO-CPU STORAGE SERVER

(57) Abstract

A storage server (110) for efficiently retrieving data from a plurality of disks (212) in response to user access requests. The server comprises a plurality of processors (302) coupled to disjoint subsets of disks, and a custom non-blocking packet switch (220) for routing data from the processors to users. By tightly coupling the processors to disks and employing an application-specific switch, congestion and disk scheduling bottlenecks are minimized. By making efficient use of bandwidth, the architecture is also capable of receiving real-time data streams from a remote source and distributing these data streams to requesting users. The architecture is particularly well suited to video-on-demand systems in which a video server stores a library of movies and users submit requests to view particular movies.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the P

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

TIGHTLY-COUPLED DISK-TO-CPU STORAGE SERVER

This application claims benefit of U.S. Provisional patent application serial number 60/127,116, filed March 5 31, 1999 and incorporated herein by reference.

The present invention relates to a storage server for retrieving data from a plurality of disks in response to user access requests. In particular, the invention relates to a multi-processing architecture in which a plurality of processors are coupled to disjoint subsets of disks, and a non-blocking cross bar switch routes data from the processors to users.

BACKGROUND OF THE DISCLOSURE

15 A storage server allows users to efficiently retrieve information from large volumes of data stored on a plurality of disks. For example, a video-on-demand server is a storage server that accepts user requests to view a particular movie from a video library, retrieves the requested program from disk, and delivers the program to the appropriate user(s). In order to provide high performance, storage servers may employ a plurality of processors connected to the disks, allowing the server to service multiple user requests simultaneously. In such multi-processor servers, processors issue commands to any of the disks, and a multi-port switch connecting the processors to the disks routes these commands to the appropriate disk. Data retrieved from disk is similarly routed back to the appropriate processor via the switch. Such servers use non-deterministic data routing channels for routing data. To facilitate accurate data retrieval, these channels require a sub-system to arbitrate conflicts that arise during data routing.

There are a number of problems, however, associated with such multi-processor servers. First, the switch

becomes a major source of latency. Since all data exchanged between the processors and disks pass through the switch and the data must be correctly routed to the appropriate destination, certain overhead processes must be accomplished to arbitrate routing conflicts and handle command and control issues. These overhead requirements cause a delay in data routing that produces data delivery latency. While it is possible to reduce such latency by reserving extra channel bandwidth, this approach dramatically increases the cost of the server. Second, the server is required to store all user requested data in a cache prior to delivery. Such a caching technique leads to poor cache efficiency wherein multiple copies of the same user data is stored in cache. These problems can significantly degrade the disk bandwidth and performance provided by the server, thereby limiting the number of users that can be supported by a given number of processors and disks. In commercial applications such as video-on-demand servers, however, it is imperative to maximize the number of users that can be supported by the server in order to achieve a reasonable cost-per-user such that the servers are economically viable.

Therefore, there is a need in the art for a multi-processor storage server that can service multiple access requests simultaneously, while avoiding the congestion, overhead, and disk scheduling bottlenecks that plague current systems.

SUMMARY OF THE INVENTION

The disadvantages associated with the prior art are overcome by a server comprising a plurality of server modules, each containing a single processor, that connect a plurality of Fibre Channel disk drive loops to a non-blocking cross bar switch such that deterministic data channels are formed connecting a user to a data source.

Each server module is responsible for outputting data at the correct time, and with the proper format for delivery to the users. A non-blocking packet switch routes the data to a proper output of the server for delivery to
5 users. Each server module supports a plurality of Fibre Channel loops. The module manages data on the disks, performs disk scheduling, services user access requests, stripes data across the disks coupled to its loop(s) and manages content introduction and migration. Since the
10 server module processors never communicate with any disks connected to other processor modules, there is no processor overhead or time wasted arbitrating for control of the Fibre Channel loops. As a result, the server can make the most efficient use of available bandwidth by
15 keeping the disks constantly busy.

The server modules transfer data read from the Fibre Channel loops to the non-blocking packet switch at the appropriate output rate. The packet switch then outputs data to a plurality of digital video modulators that
20 distribute the data to requesting users. Data requests from the users are demodulated and coupled to the switch. The switch routes the requests to the server controller which in turn routes the requests to an appropriate server module that contains the requested data. In this manner,
25 a user establishes a deterministic channel from their terminal (decoder) to the data source (disk drive) such that low latency data streaming is established.

BRIEF DESCRIPTION OF THE DRAWINGS

30 The teachings of the present invention can be readily understood by considering the following detailed description in conjunction with the accompanying drawings in which:

FIG. 1 depicts a high-level block diagram of a data retrieval system that includes a storage server incorporating the present invention;

FIG. 2 depicts a detailed block of the storage
5 server;

FIG. 3 depicts a block diagram of the CPCI chassis;

FIG. 4 depicts a block diagram of the Fibre Channel
Card;

FIG. 5 depicts a block diagram of an I/O circuit for
10 the non-blocking packet switch; and

FIG. 6 depicts a block diagram of a multiple server
system comprising the server of the present invention.

To facilitate understanding, identical reference
numerals have been used, where possible, to designate
15 identical elements that are common to the figures.

DETAILED DESCRIPTION

FIG. 1 depicts a client/server data retrieval system.
20 100 that employs a storage server 110 which accepts user
access requests from clients 120 via data paths 150.
Server 110 retrieves the requested data from disks within
the server 110 and outputs the requested data to the user
via data paths 150. Data streams from a remote source
25 (secondary storage 130) are received by the storage server
110 via data path 140. The data streams from the
secondary storage are generally stored within the storage
server for subsequent retrieval by clients 120.

In a video on demand (VOD) application, the clients
30 120 are the users' transceivers (e.g., modems that contain
video signal decoders and an associated communications
transmitter that facilitate bidirectional data
communications) and the data from the storage server is
modulated in a format (e.g., quadrature amplitude
35 modulation (QAM)) that is carried to the clients via a

hybrid-fiber-coax (HFC) network. The transceiver contains circuitry for producing data requests that are propagated to the storage server through the HFC network or some other communications channel (e.g., telephone system). In
5 such a VOD system, the remote source may be a "live feed" or an "over the air" broadcast as well as a movie archive.

FIG. 2 depicts a detailed block diagram of the storage server 110 coupled to a plurality of data modulator/demodulator circuits $222_1, 222_2, \dots, 222_n$
10 (collectively referred to as the modulator/demodulators 222). The storage server 110 comprises one or more server controllers 204, a server internal private network 206, a plurality of the server modules $208_1, 208_2, \dots, 208_n$ (collectively referred to as the server modules 208), a
15 plurality of input/output circuits 214, 218, and 216, and an non-blocking cross bar switch 220.

The server controller 204 forms an interface between the server internal private network 206 and a head end public network (HEPN) 202. The public network carries
20 command and control signaling for the storage server 110. To provide system redundancy, the server contains more than one server controller 204 (e.g., a pair of parallel controllers 204₁ and 204₂). These server controllers 204 are general purpose computers that route control
25 instructions from the public network to particular server modules that can perform the requested function, i.e., data transfer requests are addressed by the server controller 204 to the server module 208 that contains the relevant data. For example, the server controller 204
30 maintains a database that correlates content with the server modules 208 such that data migration from one server module 208 to another is easily arranged and managed. As discussed below, such content migration is important to achieving data access load balancing. Also,
35 the server controller 204 monitors loading of content into

-6-

the server modules 208 to ensure that content that is accessed often is uniformly stored across the server modules 208. Additionally, when new content is to be added to the storage server, the server controller 204 can
5 direct the content to be stored in an underutilized server module 208 to facilitate load balancing. Additional content can be added through the HEPN or via the network content input (NCI) 201. The NCI is coupled to a switch 203 that directs the content to the appropriate server
10 module 208. As further described below, the output ports of the switch 203 are coupled to the compact PCI chassis 210 within each of the server modules 208.

The server internal private (IP) network comprises a pair of redundant IP switches 206₁ and 206₂. These
15 switches route data packets (i.e., packets containing command and control instructions, and the like) from the server controller 204 to the appropriate server module 208.

Each of the server modules 208 comprise a compact PCI
20 (CPCI) chassis 210 and a plurality of fiber channel (FC) loops 224. Each of the FC loops 224 respectively comprises a disk array 212₁, 212₂, ... 212_n, and a bidirectional data path 226₁, 226₂, ... 226_n. To optimize communication bandwidth to the disk while enhancing
25 redundancy and fault tolerance, the data is striped across the disk arrays 212 in accordance with a RAID standard, e.g., RAID-5. Data is striped in a manner that facilitates efficient access to the data by each of the server modules. One such method for striping data for a
30 video-on-demand server that is known as "Carousel Serving" is disclosed in U.S. patent 5,671,377 issued September 23, 1997. Since the data is striped across all of the FC loops in a given server module, the striping is referred to as being "loop striped." Such loop striping enables
35 the server to be easily scaled to a larger size by simply

adding addition server modules and their respective FC loops. Additional data content is simply striped onto the additional disk arrays without affecting the data or operation of the other server modules 208 in the storage
5 server 110. The data accessed by the CPCI chassis 210 from the FC loops 224 is forwarded to the cross bar switch 220 via an input/output (I/O) circuit 214.

The cross bar switch 220 has a plurality of I/O ports that are each coupled to other circuits via I/O circuits
10 214, 216 and 218. The switch is designed to route packetized data (e.g., MPEG data) from any port to any other port without blocking. The I/O circuits 214 couple the cross bar switch 220 to the server modules 208, the I/O circuit 216 couples the cross bar switch to other
15 sources of input output signals, and the I/O circuits 218 couple the cross bar switch to the modulator/demodulator circuits 222. Although the I/O circuits can be tailored to interface with specific circuits, all the I/O circuits 214, 216, and 218 are generally identical. The I/O
20 circuits format the data appropriately for routing through the cross bar switch 220 without blocking. The switch 220 also contains ETHERNET circuitry 221 for coupling data to the HEPN 202. For example, user requests for data can be routed from the switch 221 to the server modules 208 via
25 the HEPN 202. As such, the I/O circuits 218 may address the user requests to the ETHERNET circuitry 221. Of course, the ETHERNET circuitry could be contained in the demodulator/ modulator circuits 222 such that the user requests could be routed directly from the demodulators to
30 the HEPN. The details of the switch 220 and its associated I/O circuits are disclosed below with respect to FIG. 5.

The modulator/demodulator circuits 222 modulate the data from I/O circuits 218 into a format that is
35 compatible with the delivery network, e.g., quadrature

amplitude modulation (QAM) for a hybrid fiber-coax (HFC) network. The modulator/demodulator circuits 222 also demodulate user commands (i.e., back channel commands) from the user. These commands have a relatively low data rate and may use modulation formats such as frequency shift key (FSK) modulation, binary phase shift key (BPSK) modulation, and the like. The demodulator circuits produce data request packets that are addressed by the I/O circuits 218 to an appropriate server module 208 such that the cross bar switch 220 routes the data request via the HEPN to a server module 208 that can implement the user's request for data.

FIG. 3 depicts a block diagram of the architecture of one of the CPCI chassis 210. The CPCI chassis 210 comprises a fibre channel card 302, a CPU card 306, an network card 304, and a CPCI passive backplane 300. The backplane 300 interconnects the cards 302, 304, and 306 with one another in a manner that is conventional to CPCI backplane construction and utilization. As such, the CPU card 306, which receives instructions from the server controller (204 in FIG. 2), controls the operation of both the FC card 302 and the input network card 304. The CPU card contains a standard microprocessor, memory circuits and various support circuits that are well known in the art for fabricating a CPU card for a CPCI chassis. The network card 304 provides a data stream from the NCI (201 in FIG. 2) that forms an alternative source of data to the disk drive array data. Furthermore, path 308 provides a high-speed connection from the cross bar switch 220 to the input network card. As such, information can be routed from the cross bar switch 220 through the network card 304 to the NCI 102 such that a communications link to a content source is provided.

The fibre channel card 302 controls access to the disk array(s) 212 that are coupled to the data paths 226

of each of the fibre channel loops 224. The card 302 directly couples data, typically video data, to and from the I/O circuits of the crossbar switch 220 such that a high speed dedicated data path is created from the array to the switch. The CPU card 306 manages the operation of the FC card 302 through a bus connection in the CPCI passive backplane 300.

More specifically, FIG. 4 depicts a block diagram of the fibre channel card 302. The fibre channel card 302 comprises a PCI interface 402, a controller 404, a synchronous dynamic random access memory (SDRAM) 410, and a pair of PCI to FC interfaces 406 and 408. The PCI interface interacts with the PCI backplane 300 in a conventional manner. The PCI interface 402 receives command and control signals from the CPU card (306 in FIG. 3) that request particular data from the disk array(s) 212. The data requests are routed to the PCI to FC interfaces 406 and/or 408. The data requests are then routed to the disk array(s) 212 and the appropriate data is retrieved. Depending upon which loop contains the data, the accessed data is routed through a PCI to FC interface 406 or 408 to the controller 404. The data (typically, video data that is compressed using the MPEG-2 compression standard to form a sequence of MPEG data packets) is buffered by the controller 404 in the SDRAM 410. The controller retrieves the MPEG data packets from the SDRAM 410 at the proper rate for each stream, produces a data routing packet containing any necessary overhead information to facilitate packet routing through the switch (220 in FIG. 2), i.e., a port routing header is appended to the MPEG data packet. The data packet is then sent to the cross bar switch 220. The controller may also perform packet processing by monitoring and setting program identification (PID) codes FIG. 5 depicts a block diagram of an I/O circuit 214, 216, or 218 for the MPEG

cross bar switch 220. The cross bar switch 220 is a multi-port switch wherein data at any port can be routed to any other port. Generally, the switch is fault tolerant by having two switches in each of the I/O circuits 214, 216, 218 to provide redundancy. One such switch is the VSC880 manufactured by Vitesse Semiconductor Corporation of Camarillo, California. This particular switch is a 16 port bi-directional, serial crosspoint switch that handles 2.0 Gb/s data rates with an aggregate data bandwidth of 32 Gb/s. The I/O circuits that cooperate with this particular switch are fabricated using model VSC 870 backplane transceivers that are also available from Vitesse. The I/O circuit, for example, circuit 214, comprises a field programmable gate array (FPGA) controller 502, cross bar switch interface 506, and buffer 508. The cross bar switch interface 506 is, for example, a VSC 870 transceiver. The buffer 508 buffers data flowing into and out of the cross bar switch. The buffer 508 may comprise two first in, first out (FIFO) memories, one for each direction of data flow. The FPGA controller 502 controls the data access through the buffer 508 and controls the cross bar switch interface 506. Additionally, the controller 502 contains a look up table (LUT) that stores routing information such as port addresses. The controller 502 monitors the buffered data and inspects the header information of each packet of data. In response to the header information and the routing information, the controller causes the buffered data to be passed through the cross bar switch interface and instructs the interface 506 regarding the routing required for the packet. The interface 506 instructs the cross bar switch as to which port on the cross bar switch 220 the data packet is to be routed.

The I/O circuits can perform certain specialized functions depending upon the component to which they are

-11-

connected. For example, the I/O circuits 218 can be programmed to validate MPEG-2 bitstreams and monitor the content of the streams to ensure that the appropriate content is being sent to the correct user. Although the foregoing embodiment of the invention "loop stripes" the data, an alternative embodiment may "system stripe" the data across all the disk array loops or a subset of loops.

FIG. 6 depicts a multiple server system 600 comprising a plurality of storage servers $110_1, 110_2, \dots, 110_n$, stores and retrieves data from a plurality of fiber channel loops. The data is routed from the server module side 214 of the switch to the modulator/demodulator side 218 of the switch. When a single server is used, all the ports on each side of the switch 220 are used to route data from the server modules 208 to the modulator/demodulators (222 in 208 FIG. 2).

To facilitate coupling a plurality of storage servers to one another and increasing the number of users that may be served data, one or more ports on each side of the switch are coupled to another server. Paths 602 couple the modulator/demodulator side 218 of switch 220 to the modulator/demodulator side 218 of switch 220, within server 110_2 . Similarly, path 604 couples the server side parts 214 to the server side 218 of switch 220_2 . In this manner, the switches of a plurality of servers are coupled to one another.

The multiple server system enables a system to be scaled upwards to serve additional users without substantial alterations to the individual servers. As such, if the switches have 8 ports on each side, the first server 110_1 , and last server 110_n , for example, use two ports on each side for inter-server data exchange and the remaining 6 ports to output data to users. The second through $n-1$ servers use four ports to communicate with adjacent servers, e.g., server 110_2 is connected to server

110, and 110₃. Note that the number of ports used to communicate between servers is defined by the desired bandwidth for the data to be transferred from server to server.

5 This arrangement of servers enables the system as a whole to supply data from any server module to any user. As such a user that is connected to server 110, can access data from server 110₃. The request for data would be routed by the HEPN to server 110₂ and the retrieved data
10 would be routed through switches 220₂ and 220₁, to the user.

 While this invention has been particularly shown and described with references to a preferred embodiment thereof, it will be understood by those skilled in the art
15 that various changes in form and details may be made therein without departing from the spirit and scope of the invention as defined by the appended claims.

-13-

What is claimed is:

1. A storage server (110) comprising
a plurality of server modules (208), each of said
5 server modules containing a processor;
a plurality of storage devices (212), each of said
storage devices coupled to exactly one of said modules;
and
a cross bar switch (220) coupled to said server
10 modules, where said server modules accept data requests
from a plurality of clients, each of said server modules
issues data retrieval commands only to the storage
devices coupled to each specific server module, and said
cross bar switch routes data from said server modules to
15 said clients requesting said data.
2. The storage server of claim 1, where said cross bar
switch also receives data from a remote source and routes
said data to said clients requesting said data.
- 20 3. The storage server of claim 1, where said plurality of
storage devices that are coupled to each of the server
modules are organized into fibre channel loops.
- 25 4. The storage server of claim 4 wherein data is striped
across the storage devices that are coupled to each of
the server modules.
5. The storage server of claim 1 wherein data stored in
30 said server modules is video data.
6. The storage server of claim 1 further comprising
input/out circuit (213) coupled to each port of said
bar switch.

-14-

7. The storage server of claim 1 wherein said data requests are routed through said cross bar switch to said server module.

5 8. The storage server of claim 1 wherein the data is striped across all the storage devices.

9. A video-on-demand server comprising:

a plurality of server modules (208), each of said
10 server modules containing a processor;

a plurality of disks (212), each of said disks coupled to exactly one of said modules, the disks form a Fibre Channel loop having video data striped across all of the disks connected to any one server module; and

15 a cross bar switch (220) coupled to said server modules, where said server modules accept data requests from a plurality of clients, each of said server modules issues data retrieval commands only to the disks coupled to each specific server module, and said cross bar switch
20 routes data from said server modules to said clients requesting said data.

10. The video-on-demand server of claim 9 wherein said data requests are routed through a communications network
25 to said server module.

11. A scaleable server comprising:

a first server (110) comprising a plurality of server modules coupled to a first cross bar switch;

30 a second server (110) comprising a plurality of server modules coupled to a second cross bar switch;

at least one data communications path coupled from said first cross bar switch to said second cross bar switch.

12. A method for providing a deterministic data channel from a data storage element to a user terminal (120) comprising the steps of:

- 5 propagating a data request from a user terminal to a storage server (110) via a communications network;
 routing the data request to a server module (208) within said storage server;
 addressing a fibre channel loop (212) containing a
10 storage device having data that fulfills the data request;
 retrieving the data to fulfill the data request; and
 routing the data from the server module through a cross bar switch (220) to the user terminal that requested the data.

15

13. The method of claim 12 wherein said step of routing the data request further comprises the step of:

- appending routing information to the data request prior to coupling the data request to the cross bar
20 switch.

14. The method of claim 12 wherein said step of routing the data further comprises the step of:

- appending routing information to the data prior to
25 coupling the data to the cross bar switch.

15. The method of claim 12 wherein data is striped across the storage devices that are coupled to each of the server modules.

30

1/6




FIG. 1

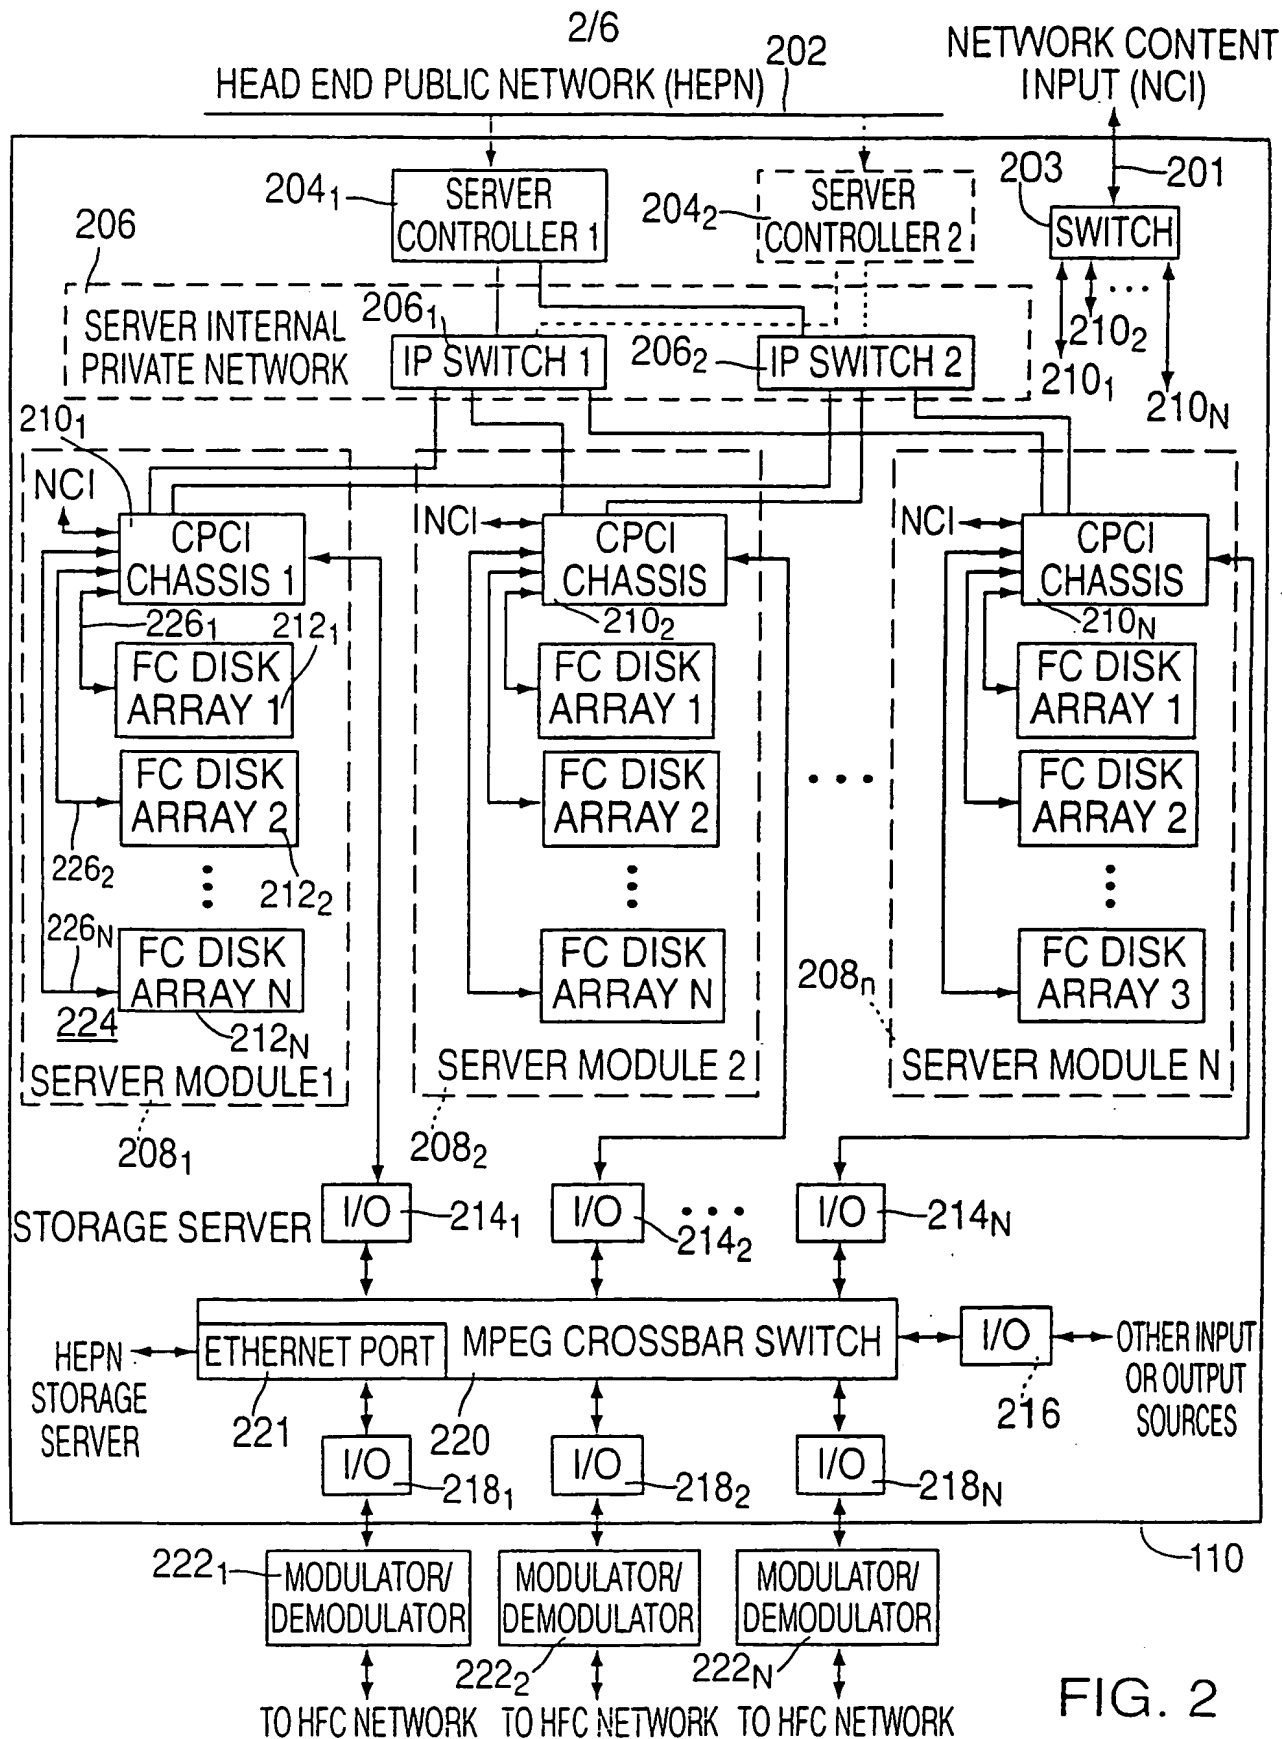


FIG. 2




FIG. 3

4/6




FIG. 4

5/6




FIG. 5

6/6




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US00/08410

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 15/16, 13/00

US CL : 709/219, 217, 203

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/219, 217, 203

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WEST DATABASE

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,862,403 A (KANAI et al) 19 January 1999, col. 6 line33 through col 9 line 10.	1-15
Y	US 5,742,599 A (LIN et al) 21 April 1998, col. 3 line 57 through col. 6 line42.	1-15
Y	US 5,694,615 A (THAPAR et al) 02 December 1997, col. 2 lin 45 through col. 4 line 28.	1-15

☐

Further documents are listed in the continuation of Box C.

☐

See patent family annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

B earlier document published on or after the international filing date

L document which may throw doubt on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

A

document member of the same patent family

Date of the actual completion of the international search

15 JUNE 2000

Date of mailing of the international search report

14 AUG 2000

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

MEHMET GECKIL

Telephone No. (703) 305-9676

Joni Hill

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☒ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER: _____**

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.